Towards a Semantics-Aware Code Transformation Toolchain for Heterogeneous Systems

Salvador Tamarit Julio Mariño
Universidad Politécnica de Madrid
Campus de Montegancedo 28660
Boadilla del Monte, Madrid, Spain
{salvador.tamarit,julio.marino}@upm.es

Guillermo Vigueras Manuel Carro
IMDEA Software Institute
Campus de Montegancedo 28223
Pozuelo de Alarcón, Madrid, Spain
{guillermo.vigueras,manuel.carro}@imdea.org

arXiv:1603.03011v2 [cs.PL] 10 Mar 2016

Abstract

Obtaining good performance when programming heterogeneous computing platforms poses significant challenges for the programmer. We present a program transformation environment, implemented in Haskell, where architecture-agnostic scientific C code with semantic annotations is transformed into functionally equivalent code better suited for a given platform. The transformation steps are formalized (and implemented) as rules which can be fired when certain syntactic and semantic conditions are met. These conditions are to be fulfilled by program properties which can be automatically inferred or, alternatively, stated as annotations in the source code. Rule selection can be guided by heuristics derived from a machine learning procedure which tries to capture how runtime characteristics (e.g., resource consumption or performance) are affected by the transformation steps.

Categories and Subject Descriptors D.3.4 [*Programming Languages*]: Processors—Code generation; C.1.3 [*Processor Architectures*]: Other Architecture Styles—Heterogeneous (hybrid) systems; I.2.2 [*Artificial Intelligence*]: Automatic Programming—Program transformation

General Terms Performance, Resource Consumption, Languages.

Keywords Rule-based Program Transformation, Semantics-aware Program Transformation, Machine Learning, High-performance computing, Heterogeneous platforms, Scientific computing, Domain-specific language, Haskell, C language.

1. Introduction

There is currently a strong trend in high-performance computing towards the integration of various types of computing elements: vector processors, GPUs being used for non-graphical purposes, FPGA modules, etc. interconnected in the same architecture. Each of these components is specially suited for some class of computations, which makes the resulting platform able to excel in performance by mapping computations to the unit best suited to execute them and is proving to be a cost-effective alternative to more traditional supercomputing architectures (Danalis et al. 2010). However, this specialization comes at the price of additional hardware and, notably, software complexity. Developers must take care of very different features to make the most of the underlying computing infrastructure. Thus, programming these systems is restricted to a few experts, which hinders its widespread adoption, increases the likelihood of bugs and greatly limits portability.

Defining programming models that ease the task of efficiently programming heterogeneous systems is an objective of many ongoing efforts, among them the European research project POLCA.[1] The project specifically targets scientific programming on heterogeneous platforms, due to the performance attained by certain hardware components for some classes of computations – e.g., GPUs and linear algebra – and to the energy savings achieved by heterogeneous computing in scientific applications characterized by high energy consumption (Danalis et al. 2010; Lindtjorn et al. 2011). Additionally, most scientific applications rely on a large base of existing algorithms that must be ported to the new architectures in a way that gets the most out of their computational strengths, while avoiding pitfalls and bottlenecks, and preserving the meaning of the original code. Porting is carried out by transforming or replacing certain fragments of code to improve their performance in a given architecture while preserving their meaning. Unfortunately, (legacy) code often does not spell its meaning or the programmer's intentions clearly, although scientific code usually follows patterns rooted in its mathematical origin.

Our goal is to obtain a framework for transformation of scientific code where the validity of a given transformation is guided by high-level annotations expressing the mathematical foundation of the source code. Despite the broad range of compilation and refactoring tools available (Bagge et al. 2003; Visser 2004; Schupp et al. 2002), no existing tool fits the needs of the project (see Section 2 for further details), so we decided to implement our own transformation framework, including a domain specific language for the definition of semantically sound code transformation rules (STML), and a transformation engine working at *abstract syntax tree* (AST) level.[2]

Therefore, we have developed a program transformation environment, implemented in Haskell, where architecture-agnostic C code is transformed into a functionally equivalent one better suited for a given platform. Haskell provides a good abstraction to represent and manipulate ASTs. Additionally, the use of pattern matching instead of visitor pattern avoids a lot of boilerplate code, making the system less error prone. Rules are written using a C-like syntax called STML, inspired by CTT (Boekhold et al. 1999) and CML (Brown et al. 2005), which makes it easy for C programmers to understand their meaning and to define them, while the rules can transparently access core functionality provided by the Haskell rewriting engine and be accessed by it to select rules and blocks of code where these rules can be safely applied. The tool is designed to use external facilities to help in selecting which rules have to be applied; we are developing a machine learning-based tools which work as external oracles to automate the selection of

[1] *Programming Large Scale Heterogeneous Infrastructures*, http://polca-project.eu.

[2] http://goo.gl/yuOFiE

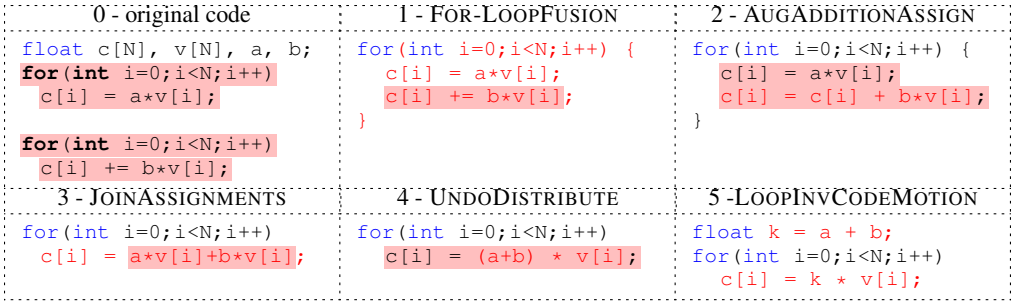

Figure 1. A sequence of transformations of a piece of C code to compute $\mathbf{c} = a\mathbf{v} + b\mathbf{v}$.

the most promising transformation chain(s) (Vigueras et al. 2016). The tool also includes an interactive mode to allow for more steering by expert users. Finally, when a optimal code is reached, it is translated to adapt to the target platform programming model.

Fig. 1 shows a sample code transformation sequence, containing the original fragment of C code along with the result of applying *loop-fusion*, reorganizing assignments, algebraic rewriting based on *distributivity*, and moving *invariant* expressions out of the body loop. Some of these transformations are currently done by existing optimizing compilers. However, they are usually performed internally, at the IR level, and without any possibility for user intervention or tailoring, which falls short to cater for many relevant situations which we want to address:

- Most compilers are designed to work without (or with minimal) human intervention, relying solely on static analysis. While when this is possible the situation is optimal, in many cases static analysis cannot uncover the underlying properties that a programmer may know. For example, in Fig. 1 any compiler would rely on native knowledge of the properties of multiplication and addition. However, if these operations were substituted by calls to implementations of operations with the same properties (distributivity, associativity, commutativity) such as operations on matrices, the transformation presented would be feasible but unlikely to be performed by a compiler relying solely on static analysis.
- Most compilers have a set of *standard* transformations which are useful in common cases for usual architectures — usually Von Neumann–based CPU architectures. However, when CPU-generic code is to be adapted for a specific architecture (e.g., FPGA, GPGPU) the transformations to be made prior to the implementation are not trivial and fall outside those usually implemented in standard compilers. Even more, compilers such as ROCCC (Jac 2012), which accept a subset of the C language and generate executables or lower-level code for a specific architecture, need the input code to follow specific coding patterns, which our tool can help generate.
- Related to the previous point, transformations to generate code to be amenable to be compiled down to some hybrid architecture can be sometimes complex and are better expressed at a higher level rather than inside a compiler's architecture. That could need users to come up with transformations which are better suited for a given coding style or application domain. Therefore, giving programmers the possibility of defining transformations at a higher level and as plugins for a compiler greatly enlarges the set of scenarios where automatic program manipulation can be applied.

Figure 2 gives an overview of the tool. It is designed to work in two stages: a *transformation* phase (Section 3) and a *translation* phase (Section 4). The former rewrites the original input code into an intermediate format, in the same input language, but which follows coding patterns closer to what would be necessary to compile into the destination architecture. The tools implementing the transformation phase use a rule library, written in a DSL (Sections 3.1 and 3.2), which defines the available transformations.

The transformed code can then be fed to compilers which accept C code adapted to the targeted architecture or be translated into code which can then be compiled for the appropriate architecture by tools which do not accept (sequential) C code.[3] Most of the work has currently been put into the transformation phase. Initial work on the translation phase shows encouraging results and point to next steps which we explain in more detail in Section 6.

2. Related Work

Rule-based program transformation is a large and fruitful area (Visser 2005; van Wijngaarden and Visser 2003). Our proposal has as distinguishing features the focus on scientific code and the aim to achieve efficiency on heterogeneous platforms, optionally using external *oracles* to guide which transformations to apply. Also, there is a crucial distinction between systems that generate new code from the mathematical model of an implementation into a new model and then generating new code, and those which use mathematical properties to transform an existing code. The former (automatic code synthesis) has long been subject of research and can generate underperforming code because of its generality. The latter usually requires that the initial code is in some "canonical" form. Our approach, based on chaining sound small-step code transformations, tries to avoid those problems.

CodeBoost (Bagge et al. 2003), built on top of Stratego-XT (Visser 2004), performs domain-specific optimizations to C++ code following an approach similar in spirit to our proposal. User-defined rules specify domain-specific optimizations; code annotations are used as preconditions and inserted as postconditions during the rewriting process. Concept-based frameworks such as Simplicissimus (Schupp et al. 2002) transform C++ based on user-provided algebraic properties. The rule application strategy can be guided by the cost of the resulting operation, although this is done at the expression level (and not at the statement level).

The transformation of C-like programs so as to optimize parallelism in its compilation into a FPGA is treated in Handel-C (Brown et al. 2005). It is however focused on a synchronous language, and therefore some of its assumptions are not valid in more general settings. A completely different approach is to use linear algebra to transform the mathematical specification of concrete scientific algorithms (Franchetti et al. 2006; Fabregat-Traver

[3] For example, and in our case, MaxJ (Maxeler Technologies 2016) code can be generated or C code with OpenMP annotations or with MPI calls.

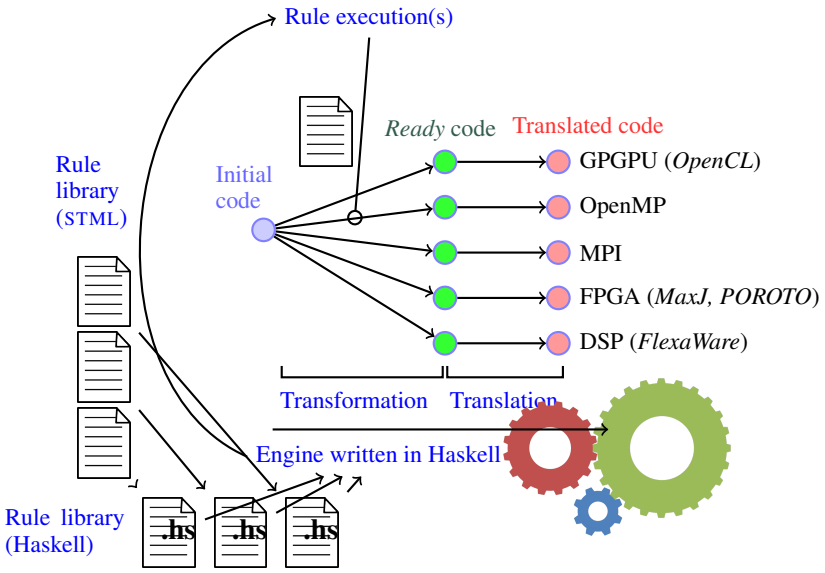

Figure 2. Anatomy of the transformation tool.

and Bientinesi 2013; Di Napoli et al. 2014). Here, the starting point is a mathematical formula and, once the formula is transformed, code is generated for the resulting expression. However, the good acceleration factors over hand-tuned code shown happens only for those algorithms, and applying the ideas to other contexts – like the aforementioned reuse of legacy code – does not seem straightforward.

There are some language-independent transformation tools that share some similarities with our approach. The most relevant are Stratego (Visser 2004), TXL (Cordy 2006), DMS (Baxter et al. 2004) and Rascal (Klint et al. 2009). Stratego is more oriented to strategies than to rewriting rules. It is not designed to add analyzers as plugins and, additionally, it does not support pragmas and does not keep a symbol table. The last two disadvantages are shared with TXL. DMS is a powerful, industrial transformation tool that is not free and there is not too much information of how it works internally, and its overall open documentation is scarce. It transforms programs by applying rules until a fix point is reached, so the rules should be carefully defined so that they do not produce loops in the rewriting stage. Finally, Rascal is still in alpha state and only available as binary, so the source code is not accessible.

3. Source-to-Source Transformations

The code transformation tool has two main components: a parser for the input language and the transformation rules written in a domain-specific language, which translates the rules into Haskell for faster execution, and an engine to perform source-to-source C code transformations based on these rules, possibly making use of information provided in code annotations given as *pragmas*.

The transformation rules are written in a domain-specific language which concisely expresses patterns which match input code and describe the skeleton of the code to generate. Rules state the conditions required to soundly apply a given transformation as well. The annotations can capture properties at two different levels: High-level properties which describe algorithmic structures and low-level properties which describe details of the procedural code. The decision of whether to apply a given transformation depends on several factors:

- First, it is necessary to ensure that applying a rule at a certain point is sound. Program properties to be matched with rule conditions can, in simple cases, be automatically verified by the parser. If this is not the case, they can be stated as annotations in the source code. These annotations may come from external tools, such as static analyzers (to extract, for example, data dependency and type information), or be provided by a programmer.
- Second, it would be necessary to ascertain whether the transformation will (eventually) improve efficiency, which is far from trivial. An interactive mode which makes it possible for a user to select which rule has to be applied is available. However, the difficulty of deciding the best rule at every moment and for every architecture, and the overhead this imposes, makes us lean towards an automatic procedure. We are developing a machine learning-based oracle (Vigueras et al. 2016) which decides, for every destination architecture, which rule to apply based on estimations of candidate transformation chains to generate code with the best possible performance.

We will next introduce the code annotations and then we will continue with the rule language. We will close this section with a description of the interaction between the transformation tool and the machine learning-based oracle.

3.1 High-Level Annotations

Information regarding the code to be transformed can be provided with a high-level interface where annotations describe semantic features of the code following a functional programming style. This makes it possible to capture algorithmic skeletons at a higher level of abstraction and to express properties of the underlying code. For instance, `for` loops expressing a mapping between an input and an output array can be annotated directly with a `map` pragma such as `#pragma polca map F v w`. This would indicate that the loop traverses the input array `v`, applies function `F` to each element in `v`, and stores the result in `w`. Additionally, and for the annotation to be correct, we require that `F` is pure, that `v` and `w` have the same length, and that every element in `w` is computed only from the corresponding element in `v`.

POLCA ANNOTATION
`#pragma polca map F v w`
TRANSLATION TO STML

`#pragma stml reads v in {0}`
`#pragma stml writes w in {0}`
`#pragma stml same_length v w`
`#pragma stml pure F`
`#pragma stml iteration_space 0 length(v)`
`#pragma stml iteration_independent`

POLCA ANNOTATION
`#pragma polca fold F INI v a`
TRANSLATION TO STML

`#pragma stml reads v in {0}`
`#pragma stml reads output(INI)`
`#pragma stml writes a`
`#pragma stml pure F`
`#pragma stml iteration_space 0 length(v)`

POLCA ANNOTATION
`#pragma polca itn F INI n w`
TRANSLATION TO STML

`#pragma stml reads output(INI)`
`#pragma stml reads n`
`#pragma stml writes w`
`#pragma stml pure F`
`#pragma stml iteration_space 0 n`

POLCA ANNOTATION
`#pragma polca zipWith F v w z`
TRANSLATION TO STML

`#pragma stml reads v in {0}`
`#pragma stml reads w in {0}`
`#pragma stml writes z in {0}`
`#pragma stml same_length v w`
`#pragma stml same_length v z`
`#pragma stml pure F`
`#pragma stml iteration_space 0 length(v)`
`#pragma stml iteration_independent`

POLCA ANNOTATION
`#pragma polca scanl F INI v w`
TRANSLATION TO STML

`#pragma stml reads output(INI)`
`#pragma stml reads v in {0}`
`#pragma stml reads w in {0}`
`#pragma stml writes w in {1}`
`#pragma stml pure F`
`#pragma stml iteration_space 0 length(v)`

Table 1. POLCA annotations and its translation into STML annotations.

The boxes labeled "POLCA ANNOTATION" in Table 1 list the main high-level annotations that can be currently used. For illustrative purposes, Listing 1 shows an annotated version of the code in Figure 1. The annotation in the second loop uses the constructor `zip` to express that a pair of arrays is treated as an array of pairs —that is necessary due to the signature of `map`.

Listing 1. Annotations for the initial code in Figure 1.

```
float c[N], v[N], a, b;

#pragma polca map BODY1 v c
for(int i=0;i<N;i++)
#pragma polca def BODY1
#pragma polca input v[i]
#pragma polca output c[i]
  c[i] = a*v[i];
```

Listing 2. BNF grammar for STML annotations.

```
<code_prop_list> ::= "#pragma stml" <code_prop>
   | "#pragma stml" <code_prop> <code_prop_list>

<code_prop>  ::= <exp_prop> <exp>
   | [<op>] <op_prop> <op>
   | <mem_access> <exp> ["in" <offset_list>]
   | "write("<exp>") =" <location_list>
   | "same_length" <exp> <exp>
   | "output("<exp>")" | <loop_prop>

<loop_prop>  ::= "iteration_independent"
   | "iteration_space" <parameter> <parameter>

<exp_prop>   ::= "appears" | "pure" | "is_identity"

<op_prop>    ::= "commutative" | "associative"
   | "distributes_over"

<mem_access> ::= "writes" | "reads" | "rw"

<location_list> ::=
   "{" <c_location> {"," <c_location>} "}"

<offset_list> ::= "{" <INT> {"," <INT>} "}"

<exp>          ::= <C_EXP> | <C_VAR> | <polca_var_id>

<op>           ::= <C_OP> | <C_VAR> | <polca_var_id>

<c_location>   ::= <C_VAR> | <C_VAR>("["<C_EXP>"]")+

<parameter>    ::= <c_location> | <polca_var_id> | <INT>
```

```
#pragma polca map BODY2 zip(v,c) c
for(int i=0;i<N;i++)
#pragma polca def BODY2
#pragma polca input (v[i],c[i])
#pragma polca output c[i]
  c[i] += b*v[i];
```

3.2 STML Properties

A source-to-source transformation tool for a procedural language requires information about the code (i.e. code properties) to ensure that transformations are correctly applied. These properties are of a lower level than those directly captured by high-level annotations, because they have to deal with characteristics pertaining to a concrete programming language. For example, a purely functional semantics for the high-level annotations cannot naturally capture some aspects of imperative languages such as destructive assignment or aliasing, since it would have to follow the referential transparency common in functional languages. Likewise, manual memory management (explicit memory allocation and deallocation) would be difficult to represent. In our framework, these procedural code-level characteristics are expressed in a language we have termed STML (*Semantic Transformation Meta-Language*) which is used both in the code annotations and in the conditions of the transformation rules.

3.2.1 Syntax and Semantics of STML Annotations

The BNF grammar for STML annotations is shown in Listing 2. An intuitive explanation of its semantics follows:

- `<code_prop>` refers to the different code properties expressed through STML annotations. By means of these annotations the user can provide information about read/write accesses to variables, loop properties, etc. The different options for this term are described below.

- `[<exp>] <exp_prop> <exp>`: The term `<exp_prop>` indicates properties about code expressions of the statement that is immediately below the annotation. Some examples are:
  - `appears <exp>`: There is at least one occurrence of `<exp>` in the statement below.
  - `pure <exp>`: Expression `<exp>` is pure, i.e. it neither has side effects nor writes on any memory location.
  - `is_identity <exp>`: Expression `<exp>` is the identity element. This annotation must be preceded by high-level annotations that define the group or field in which `<exp>` is the identity element.
- `[<op>] <op_prop> <op>`: The term `<op_prop>` can be an unary or binary operator indicating properties about operators. Some examples are:
  - `commutative <op>`: Operator `<op>` has the commutative property: $\forall x, y.\ f(x, y) = f(y, x)$
  - `associative <op>`: Operator `<op>` has the associative property: $\forall x, y, z.\ f(f(x, y), z) = f(x, f(y, z))$
  - `<op> distributes_over <op>`: The first operator distributes over the second operator: $\forall x, y, z.\ g(f(x, y), z) = f(g(x, z), g(y, z))$
- `"write("<exp>")="<location_list>`: This term states the list of memory locations written on by expression `<exp>`, where `<location_list>` is a list of variables (scalar or array type) in the C code. For example, `write(c = a + 3)= {c}` and `write(c[i++] = a + 3)= {c[i], i}`
- `<mem_access> <exp> ["in"<offset_list>]`: The term `<mem_access>` states properties about the memory accesses to expression `<exp>` of the statement (or statements) that immediately follow. When the expression `<exp>` is an array, the list of array positions accessed can be referenced through `"in"<offset_list>` where `<offset_list>` is the list of array positions accessed either for writing or reading, depending on the access stated by `<mem_access>`. Some examples are:
  - `writes <exp>`: This annotation specifies that the set of statements associated to the STML annotation writes into location identified by `<exp>`.
  - `writes <exp> "in"<offset_list>`: This annotation is similar to the previous one, but for non-scalar variables within loops. It specifies that for each `i`-th iteration of the loop, an array identified by `<exp>` is written to in the locations whose offset with respect to the index of the loop is contained in `<offset_list>`. For example:

```
#pragma stml writes c in {0}
for (i = 0; i < N; i++)
   c[i] = i*2;

#pragma stml writes c in {-1,0}
for (i = 1; i < N; i++){
   c[i-1] = i;
   c[i]  = c[i-1] * 2;
}
```

  - `reads <exp>`: The set of statements associated to the STML annotation read from location `<exp>`.
  - `reads <exp> "in"<offset_list>`: This annotation is similar to the previous one, but for non-scalar variables within loops: for the `i`-th iteration of the loop, the array identified by `<exp>` is accessed to read locations with offset values contained in `<offset_list>`. An example follows:

```
#pragma stml reads c in {0}
for (i = 0; i < N; i++)
   a += c[i];

#pragma stml reads c in {-1,0,+1}
for (i = 0; i < N; i++)
   a += c[i-1]+c[i+1]-2*c[i];
```

  - `rw <exp>` The set of statements associated to the STML annotation reads and writes from / into location `<exp>`.
  - `rw <exp> "in"<offset_list>` This annotation is similar to the previous one, but for non-scalar variables within loops. The annotation specifies that for the `i`-th iteration of the loop, an array identified by `<exp>` is accessed to read and/or write locations with offset values contained in `<offset_list>`.
- `<loop_prop>`: This term represents annotations related with loop properties
  - `"iteration_space"<parameter> <parameter>` This annotation states the iteration space limits of the `for` loop associated with the annotation. An example would be:

```
#pragma stml iteration_space 0 N-1
for (i = 0; i < N; i++)
   c[i] = i*2;
```

  - `"iteration_independent"` This annotation is used to state that there is no loop-carried dependencies in the body of the loop associated to this annotation.
- `"same_length"<exp> <exp>`: This annotation states that two arrays in the C code, given as parameters in `<exp>`, have the same length.
- `"output("<exp>")"`: This annotation is used to reference the output of a block of code identified by `<exp>`.

3.2.2 Translation from High-Level to STML Annotations

When procedural code is decorated with high-level annotations, the annotated code is assumed to implement the computation expressed in the annotation. Additionally, the elements which appear in the annotation are supposed to follow a functional semantics, such as referential transparency. Using this interpretation, lower-level STML properties can be inferred for the annotated code and used to decide which transformations are applicable.

For example, consider a loop annotated with a `map F v w`, like the first one in Listing 1, where `F` is `BODY1`, `v` is `v`, and `w` is `c`. In this context, we assume that:

- `F` behaves as if it had no side effects. It may read and write from/to a global variable, but it should behave as if this variable did not implement a state for `F`. For example, it may always write to a global variable and then read from it, and the behavior of other code should not depend on the contents of this variabe.
- `v` and `w` are arrays of the same size.
- For every element of `w`, the element in the i-th position is computed by applying `F` to the element in the i-th position of `v`.
- The applications of `F` are not assumed to be done in any particular order (i.e., they can go from `v[0]` upwards to

Listing 3. Code from Listing 1 after translating high-level annotations to STML.

```
float c[N], v[N], a, b;

#pragma polca map BODY1 v c
#pragma stml reads v in {0}
#pragma stml writes c in {0}
#pragma stml same_length v c
#pragma stml pure BODY1
#pragma stml iteration_space 0 length(v)
#pragma stml iteration_independent
for(int i = 0; i < N; i++)
#pragma polca def BODY1
  c[i] = a*v[i];

#pragma polca map BODY2 zip(v,c) c
#pragma stml reads (v in {0}, c in {0})
#pragma stml writes c in {0}
#pragma stml same_length zip(v,c) c
#pragma stml pure BODY2
#pragma stml iteration_space 0 length(zip(v,c))
#pragma stml iteration_independent
for(int i = 0; i < N; i++)
#pragma polca def BODY2
  c[i] += b*v[i];
```

`v[length(v)-1]` or in the opposite direction. Therefore all applications of `F` should be independent from each other.

The STML properties inferred from some high-level annotations are shown in Table 1. If we focus again our attention on the translation of `map`, the STML annotations mean that:

- Iteration i-th reads from `v` in the position i-th (it actually reads in the set of positions $\{i+0\text{-th}\}$, since the set of offsets it reads from is $\{0\}$).
- Iteration i-th ultimately writes on `w` in the position i-th (same comment as before).
- `v` and `w` have the same length.
- `F` behaves as if it did not have side effects.
- `F` is applied to `v` and `w` in the indexes ranging from 0 to $length(\mathtt{v})$.

The correspondence of other high-level annotations (`foldl`, `itn`, `zipWith` and `scanl`) into STML annotations is shown in Table 1. An extended explantation of these annotations can be found in (Rubio et al. 2015). As in the case of the `map` annotation, the derived STML assertions provide properties of the procedural code which are used to ensure that transformation rules are correctly applied. An example is shown in Listing 3, which corresponds to the translation into STML annotations of the code in Listing 1, which implements the computation of a `map`.

The STML annotations for the code are internally used by the source-to-source transformation tool to decide which transformations can be applied. For example in this case, the FORLOOPFUSION transformation (Table 3) needs certain non-syntactical properties to be met. A description of the STML properties used for the transformations in Figure 1 is shown in Table 2.

3.2.3 Using External Tools

Besides the properties provided by the user, either written directly in STML or deduced from high-level annotations, additional properties can be obtained from external tools. This is useful as it would relieve users from having to write a large number of annotations which state many low-level details. These properties can be made available to the transformation tool by writing them as STML annotations.

We are currently using Cetus (Dave et al. 2009) to automatically produce STML annotations. Cetus is a compiler framework, written in Java, to implement source-to-source transformations. We have modified it (which is allowed by its license) to add some new analyses and to output the properties it infers as STML pragmas annotating the input code.

If the annotations automatically inferred by external tools contradict those provided by the user, the properties provided by the user are preferred to those deduced from external tools, but a warning is issued nonetheless.

```
join_assignments {
   pattern: {
      cstmts(s1);
      cexpr(l) = cexpr(e1);
      cstmts(s2);
      cexpr(l) = cexpr(e2);
      cstmts(s3);
   }
   condition: {
      no_write(cstmts(s2),{cexpr(l),cexpr(e1)});
      no_read(cstmts(s2),{cexpr(l)});
      pure(cexpr(e1));
      pure(cexpr(l));
   }
   generate: {
      cstmts(s1);
      cstmts(s2);
      cexpr(l) = subs(cexpr(e2),cexpr(l), cexpr(e1));
      cstmts(s3);
   }
}
```

Figure 3. STML rule for JOINASSIGNMENTS.

3.3 STML Rules

As an example of a rewriting rule written in STML, let us consider the code for the rule JOINASSIGNMENTS shown in Figure 3. The rule is applied when the program / library that we are transforming has a piece of code matching the pattern in the `pattern` section and it fulfills the conditions in the `condition` section. The resulting code corresponds to the pattern in the `generate` section, where expressions matched in the `pattern` are replaced in the `generate`d code. In this case one assignment is removed by propagating the expression in the assignment's right hand side.

STML uses metavariables to match components of the initial code. These metavariables are *tagged* to denote explicitly which kind of component they should match. For example, a metavariable `v` can be tagged `cexpr(v)` to denote that it can only match an expression, `cstmt(v)` to denote that it can only match a statement, or `cstmts(v)` to denote that it can match a sequence of statements. In the example in Figure 3, `s1`, `s2` and `s3` should be (sequences of) statements, and `e1`, `e2` and `l` are expressions.

There are additional conditions and primitives to help in generating code; these are presented and briefly described in Tables 5 and 4, respectively. In these tables, `E` represents an expression, `S` represents a statement, while `[S]` represents a sequence of statements. Additionally, the function `bin_oper(`$\mathtt{E_{op}}$`,`$\mathtt{E_l}$`,`$\mathtt{E_r}$`)` matches or generates a binary operation ($\mathtt{E_l}$ $\mathtt{E_{op}}$ $\mathtt{E_r}$) and can be used in the sections `pattern` and `generate`. The section `generate` can also state, using `#pragmas`, new properties which hold in the resulting code.

$s \not\mapsto l$	statements s do not write into location l: $l \notin \text{writes}(s)$
$s \not\leftarrow l$	statements s do not read the value in location l
$s_1 \not\mapsto s_2$	statements s_1 do not write into any location read by s_2
$s_1 \not\leftarrow s_2$	statements s_1 do not read from any location written by s_2
$s_1 \underset{-a[l]}{\not\mapsto} s_2$	same predicate than previous one but not taking into account locations referred through arrays.
$s_1 \underset{a[l]}{\overset{<}{\not\mapsto}} s_2$	statements s_1 do not write into any previous location corresponding to an index array read by s_2
e pure	expression e is *pure*, i.e. does not have side effects nor writes any memory locations.
writes(s)	set of locations written by statements s.
g distributes_over f	$\forall x, y, z.\ g(f(x, y), z) \approx f(g(x, z), g(y, z))$
l fresh	l is the location of a *fresh* identifier, i.e. does not clash with existing identifiers if introduced in a given program state.

Table 2. Predicates used to express conditions for the application transformation rules in Table 3.

$$\texttt{for}\,(l{=}e_{ini}\texttt{;}\,rel(l, e_{end})\texttt{;}\,mod(l))\ \{s_1\}\ \texttt{for}\,(l{=}e_{ini}\texttt{;}\,rel(l, e_{end})\texttt{;}\,mod(l))\ \{s_2\} \Rightarrow \texttt{for}\,(l{=}e_{ini}\texttt{;}\,rel(l, e_{end})\texttt{;}\,mod(l))\ \{s_1\,\texttt{;}\,s_2\}$$

when rel pure, $(s_1\,\texttt{;}\,s_2) \not\mapsto \{l, e_{ini}, e_{end}\}$, writes$(mod(l)) \subseteq \{l\}$, $s_1 \underset{-a[l]}{\not\mapsto} s_2, s_2 \underset{-a[l]}{\not\mapsto} s_1, s_2 \underset{a[l]}{\overset{<}{\not\mapsto}} s_1$

(FOR-LOOPFUSION)

$l\ \texttt{+=}\ e\texttt{;} \Rightarrow l\ \texttt{=}\ l\ \texttt{+}\ e\texttt{;}$
when l pure

(AUGADDITIONASSIGN)

$s_1\texttt{;}\ l\ \texttt{=}\ e_1\texttt{;}\ s_2\texttt{;}\ l\ \texttt{=}\ e_2\texttt{;}\ s_3\texttt{;} \Rightarrow s_1\texttt{;}\ s_2\texttt{;}\ l\ \texttt{=}\ e_2[e_1/l]\texttt{;}\ s_3\texttt{;}$
when l, e_1 pure, $s_2 \not\mapsto \{l, e_1\}, s_2 \not\leftarrow l, s_2 \not\mapsto e_1$

(JOINASSIGNMENTS)

$f(g(e_1, e_3), g(e_2, e_3)) \Rightarrow g(f(e_1, e_2), e_3)$
when e_1, e_2, e_3 pure, g distributes_over f

(UNDODISTRIBUTE)

$\texttt{for}\ (e_1\texttt{;}e_2\texttt{;}e_3)\ \{s_b\} \Rightarrow l\ \texttt{=}\ e_{inv}\texttt{;}\ \texttt{for}\ (e_1\texttt{;}e_2\texttt{;}e_3)\ \{s_b[l/e_{inv}]\}$
when l fresh, e_{inv} occurs in s_b, e_{inv} pure, $\{s_b, e_3, e_2\} \not\mapsto e_{inv}$

(LOOPINVCODEMOTION)

Table 3. Source code transformations used in the example of Figure 1.

Function/Construction	Description
`subs((S|[S]|E),`$\texttt{E}_\texttt{f}$`,`$\texttt{E}_\texttt{t}$`)`	Replace each occurrence of $\texttt{E}_\texttt{f}$ in `(S|[S]|E)` for $\texttt{E}_\texttt{t}$.
`if_then:{`$\texttt{E}_\texttt{cond}$`; (S|[S]|E);}`	If $\texttt{E}_\texttt{cond}$ is true, then generate `(S|[S]|E)`.
`if_then_else:{`$\texttt{E}_\texttt{cond}$`; (S|[S]|E)`$_\texttt{t}$`; (S|[S]|E)`$_\texttt{e}$`;}`	If $\texttt{E}_\texttt{cond}$ is true, then generate `(S|[S]|E)`$_\texttt{t}$ else generate `(S|[S]|E)`$_\texttt{e}$.
`gen_list:{[(S|[S]|E)];}`	Each element in `[(S|[S]|E)]` produces a different rule consequent.

Table 4. Rule language constructions and functions for `generate` rule section.

3.4 Rule Selection

Our code transformation tool performs chaining of rule applications which go from initial code to a final form where some conditions are met —see, e.g., Figure 1. In general, several rules can be applied at multiple code points in the rewriting process. Deciding whether a given rule is applicable or not depends on whether rule conditions are met or not. In practice, it is often the case that the tool does not have information enough to decide whether a property holds or not, and therefore it cannot be decided whether a rule which requires that property can be applied or not. Thus, the system distinguishes between definitely "applicable", "definitely not applicable" and "probably (not) applicable" transformation steps. For all applicable rules, deciding which one should be chosen should be based on whether that rule can contribute to an eventual improvement of the performance of the final code with respect to the original one. We provide two ways to perform rule selection: the possibility of having interaction with a human and a generic interface to communicate with external tools.

3.4.1 Interactive Rule Selection

Our tool provides an interface to make it possible an interactive transformation: the user is presented with the possible rules to be applied to some code and the code before and after applying the rule. This can be useful to refine rules or to perform refactoring with general purposes (e.g., improve readability, make changes which impact large sections of code, perform repetitive but delicate main-

Function	Description
`no_write((S|[S]|E)1, (S|[S]|E)2)`	True if `(S|[S]|E)1` does not write in any location read by `(S|[S]|E)2`.
`no_write_except_arrays ((S|[S]|E)1,(S|[S]|E)2,E)`	Same than previous, but not taking into arrays accessed using `E`.
`no_write_prev_arrays ((S|[S]|E)1 (S|[S]|E)2, E)`	True if all array writes indexed using `E` in `(S|[S]|E)1` do not access previous locations to array reads indexed using `E` in `(S|[S]|E)2`.
`no_read((S|[S]|E)1, (S|[S]|E)2)`	True if `(S|[S]|E)1` does not read in any location written by `(S|[S]|E)2`.
`pure((S|[S]|E))`	True if `(S|[S]|E)` does not write in any location.
`writes((S|[S]|E))`	Locations written by `(S|[S]|E)`.
`distributes_over(E1,E2)`	True if operation `E1` distributes over operation `E2`.
`occurs_in(E,(S|[S]|E))`	True if expression `E` occurs in `(S|[S]|E)`.
`fresh_var(E)`	Indicates that `E` should be fresh.
`is_identity(E)`	True if `E` is the identity.
`is_assignment(E)`	True if `E` is an assignment.
`is_subseteq(E1,E2)`	True if $E_1 \subseteq E_2$

Table 5. Rule language functions for the section `condition` of a rule.

tenance, ...) which may not be related to improving performance or to adapt existing code to a given platform. Auxiliary programs can be used to show the differences in the code before and after applying a given transformation step (for example, Meld (Willadsen 2016)) that can help the user in this process.

3.4.2 Oracle-Based Rule Selection

Due to its fine-grained nature, manual rule selection is in general not scalable: in our experience, it is not a realistic possibility for even small programs with a reduced set of rules. Therefore, mechanizing as much as possible this process is a must. A straightforward possibility is to select at each step the rule[4] which reduces some metric. This may however make the transformation be trapped in local minima, and in our experience, in many cases it is necessary to apply transformations which temporarily reduce the quality of the code because they enable the application of further transformations.

A possibility to work around this problem is to explore a bounded number of possible rule applications "in the vicinity of" some state. How many steps should be taken in this exploration is something which needs to be decided, since taking too few steps would not make it possible to leave a local minima. Given that in our experience the number of rules that can be applied in most states is high (typically in the order of the tens), increasing the diameter of the boundary to be explored can cause an exponential explosion in the number of states to be evaluated, even for such a bounded search.

Additionally, and since rules are a parameter for the transformation engine, the user can introduce rules which are the inverse of each other (which may lead to infinite loops) or which duplicate the code, such as loop unfolding rules, making it possible to apply repeatedly the same rule to states which are different (e.g., to different code configurations).

Among the possibilities to select a rule which is part of a promising sequence which leads to good code for a given platform (see Section 4), we are exploring the use of machine learning techniques based on *reinforcement learning* (Vigueras et al. 2016). From the point of view of the transformation engine, the selection tool works as an *oracle* which, given a code configuration and a set of applicable rules, returns which rule should be applied. We will describe now an abstract interface to an external rule selector, which can be applied not only to the current oracle, but to other similar tools.

[4] Or one of them, should there be several candidates.

The interface of the transformation tool is described by functions

- $AppRules(Code) \rightarrow \{(Rule, Pos)\}$
- $Trans(Code_i, Rule, Pos) \rightarrow Code_o$

The function $AppRules$ determines the possible transformations applicable to a given code. It returns, for a given input $Code$, a set of tuples containing each of a rule name $Rule$ and the code position Pos where it can be applied (e.g., the identifier of a node in the AST). On the other hand, function $Trans$ is basically the application of a given transformation step. For an input code $Code_i$, a rule name $Rule$, and a code position Pos, it returns the resulting code $Code_o$ after applying the transformation.

The interface of the reinforcement learning tool includes the functionality described at the beginning of this section – rule selection and stop condition:

- $SelectRule(\{(Code_i, \{Rule_i\})\}) \rightarrow (Code_o, Rule_o)$
- $IsFinal(Code) \rightarrow Boolean$

The function $SelectRule$ selects which code has to be transformed and which rule has to be applied to perform the transformation. Given a set of tuples containing a code $Code_i$ and a set of rules $\{Rule_i\}$, it returns a tuple containing the chosen code $Code_o$ and the rule $Rule_o$ that should be applied to $Code_o$. Likewise, function $IsFinal$ is used to know whether a given code $Code$ is considered ready for translation or not.

We are now in a position to introduce the function that defines the interaction between the transformation and the external oracle:

- **Header**: $NewCode(Code_i, \{Rule_i\}) \rightarrow (Code_o, Rule_o)$
- **Definition**: $NewCode(c, rules) =$
  $SelectRule($
  $\{(c', \{r' \mid (r', _) \in AppRules(c')\})$
  $\mid c' \in \{ Trans(c, r, p)$
  $\mid (r, p) \in AppRules(c), r \in rules\} \})$

The function $NewCode$ receives an initial code $Code_i$ and a set of rules $\{Rule_i\}$ which are candidates to be applied. It returns the code $Code_o$ resulting from the application of one of the rules from $\{Rule_i\}$ to $Code_i$. Additionally, it returns the rule $Rule_o$ that should be applied in the next transformation step, i.e., the next time $NewCode$ is invoked with $Code_o$. The rationale behind this design is that the first invocation receives all the applicable rules as

candidates to be applied, but after this first application we always have a single rule in the set $\{Rule_i\}$. We do it repeatedly until the transformation generates a code for which function $IsFinal$ returns true:

- **Complete derivation**:
  $NewCode(c_0, AllRules) \rightarrow^* (c_n, r_n)$
  where $IsFinal(c_n)$ and $\forall\ 0 < i < n.$
  $(c_i, r_i) = NewCode(c_{i-1}, \{r_{i-1}\})$ with $\neg IsFinal(c_i)$

This approach makes it unnecessary for the external oracle to have to consider positions where a transformation can be applied, since that choice is implicit in the selection of a candidate code between all possible code versions obtained using a single input rule. Furthermore, by selecting the next rule to be applied, it takes the control of the next step of the transformation. The key here is the function $SelectRule$, used in function $NewCode$. Given an input code $Code_i$ and a rule $Rule_i$,[5] $SelectRule$ selects a resulting code between all the codes that can be generated from $Code_i$ using $Rule_i$. The size of the set received by function $SelectRule$ corresponds to the total number of positions where $Rule_i$ can be applied. In this way, $SelectRule$ is implicitelly selecting a position.

4. Producing Code for Heterogeneous Systems

In the second phase of the tool (Figure 1), code for a given platform is produced starting from the result of the transformation process. The destination platform of a fragment of code can be specified using annotations like this:

```
#pragma polca mpi
```

This information is relevant, on one hand, to hint at what transformations should be applied and also to decide when the code is ready for translation. In fact, the decisions taken by the machine learning-based tool we mentioned before is partly directed by the destination architecture.

The translation to a final code for a given architecture is in most cases straightforward as it needs only to introduce the "idioms" necessary for the architecture or to perform a syntactical translation. As a consequence, there is no search or decision process: for each input code given to the translation, there is only one output code which is obtained via predefined transformations or glue code injection.

Some of the translations need specific information: for instance, knowing if a statement is performing I/O is necessary when translating to MPI, because executing this operation might need to be done in a single thread. It is often the case that this can be deduced by syntactical inspection, but in other cases (e.g., if the operation is part of a library function) it may need explicit annotations.

5. Implementation Notes

The transformation phase, which obtains C code that could be easily translated into the source language for the destination platform, is a key part of the tool. As a large part of the system was experimental (including the definition of the language, the properties, the generation of the final code, and the search / rule selection procedures), we needed a flexible and expressive implementation platform.

We initially tested well-known infrastructures such as Clang / LLVM. While they are very well supported and tested, we found that they understandably were primarily designed for compilation instead of for source-to-source program transformation, which is our main goal. When implementing complex source-to-source program transformation routines in Clang, we found that the interface offered was not really designed to perform AST transformations, and that the design documents warned that the interface could not be assumed to be stable. Additionally, the methods to transform source code had to be coded in C++ which made them verbose and full of low-level details which we did not want to deal with. Compiling rules to C++ was of course an option, but even this compilation was not going to be easy, due to the conceptual distance between the nature of the rules and the API for code manipulation, and it was dependant upon an unstable interface. Even in that case, the whole Clang would have to be recompiled after introducing new rules, which made project development and testing cumbersome, and would make the addition user-defined rules complicated.

Therefore we decided to switch to a declarative language and implement the tool in Haskell. Parsing the input code is done by means of the `Language.C` (Huber 2014) library, which returns the AST as a data structure which is easy to manipulate. In particular, we used the Haskell facilities to deal with generic data structures through the *Scrap Your Boilerplate* (SYB) library (Lammel et al. 2009). This allows us to easily extract information from the AST or modify it with a generic traversal of the whole structure.

The rules themselves are written in a subset of C, and can therefore be also parsed using `Language.C`. After reading them in, they are automatically compiled into Haskell code (contained in the file `Rules.hs` —see Figure 2) which performs the traversal and (when applicable) the transformation of the AST. This module is loaded together with the rest of the tool, therefore avoiding the extra overhead of interpreting the rules. The declarative nature of Haskell and facilities such as completely automatic memory management makes this compilation much easier than it would be in the case of compiling into C.

When it comes to rule compilation, STML rules can be divided into two classes: those which operate at the expression level (which are easier to implement) and those which can manipulate both expressions and (sequences of) statements. In the latter case, sequences of statements (`cstmts`) of an unknown size have to be considered: for example, in Figure 3, `s1`, `s2`, and `s3` can be sequences of any number of statements (including the empty sequence), and the rule has to try all the possibilities to determine if there is a match which meets the rule conditions. For this, Haskell code that explicitly performs an AST traversal needs to be generated. In the case of expressions, they are syntactically bound, and the translation of the rule is much easier.

When generating Haskell code, the rule sections (`pattern`, `condition`, `generate`, `assert`) generate the corresponding LHS's, guards, and RHS's of a Haskell function. If the conditions to apply a rule are met, the result is returned in a triplet `(rule_name, old_code, new_code)` where the two last components are, respectively, the matched and transformed sections of the AST. Note that `new_code` may contain new properties if the `generate` section of the rule defines them.

Since several rules can be applied at several locations of the AST, every rewriting step can actually return a list of tuples — one for each rule and location where that rule can be applied. As mentioned elsewhere (Section 3.4.2), besides the possibility of interacting with a user, we are studying the usage of an external oracle which determines the best candidate to apply in the next step. The transformation halts when either no more rules are applicable or when a stop condition is found, according to the oracle.

The tool is divided into four main modules:

- **`Main.hs`** implements the main workflow of the tool: it calls the parser on the input C code to build the AST, links the prag-

[5] Note that the corresponding parameter is a set only for the initial call.

mas to the AST, executes the transformation sequence (interactive or automatically) and outputs the transformed code.

- **`PragmaLib.hs`** reads pragmas and links them to their corresponding node in the AST. It also restores or injects pragmas in the transformed code.
- **`Rul2Has.hs`** translates STML rules (stored in an external file) into Haskell functions which actually perform the AST manipulation. It also reads and loads STML rules as an AST and generates the corresponding Haskell code in the **`Rules.hs`** file.
- **`RulesLib.hs`** contains supporting code used by `Rules.hs` to identify whether some STML rule is or not applicable (e.g., there is matching code, the preconditions hold, etc.) and to execute the implementation of the rule (including AST traversal, transformation, ...).

6. Conclusion

We have presented a transformation toolchain that uses semantic information, in the form of user- or machine-provided annotations, to produce code for different platforms. It has two clearly separated phases: a source-to-source transformation which generates code with the style appropriate for the destination architecture and a translation from that code to the one used in the specific platform.

We have focused until now in the initial phase, which included the specification of a DSL (STML) to define rules and code properties, a translator from this language into Haskell, a complete engine to work with these rules, and an interface to interact with external oracles (such as a reinforcement learning tool which we are developing) to guide the transformation.

The translation phase is still in an preliminary stage. However, and while it is able to translate some input code, it needs to be improved in order to support a wider range of input code. We have compared, using several metrics, the code obtained using our tool and the corresponding initial code and the results are encouraging.

As future work, we plan to improve the usability of the STML language. At the same time, we are modifying Cetus to automatically obtain more advanced / specific properties, and we are integrating profiling techniques in the process to make it easier to evaluate the whole transformation system and give feedback on it.

Simultaneously, we are investigating other analysis tools which can be used to derive more precise properties. Many of these properties are related to data dependencies and pointer behavior. We are considering, on one hand, tools like PLuTo (Bondhugula et al. 2008) and PET (Verdoolaege and Grosser 2012) (two polytope model-based analysis tools) or the dependency analyzers for the Clang / LLVM compiler. However, since they fall short to derive dependencies (e.g., alias analysis) in code with pointers, we are also considering tools based on separation logic (O'Hearn et al. 2001; Reynolds 2002) such as VeriFast (Jacobs and Piessens 2008; Jacobs et al. 2011) which can reason on dynamically-allocated, mutable structures.

Acknowledgments

Work partially funded by EU FP7-ICT-2013.3.4 project 610686 POLCA, Comunidad de Madrid project S2013/ICE-2731 N-Greens Software, and MINECO Projects TIN2012-39391-C04-03 / TIN2012-39391-C04-04 (StrongSoft), TIN2013-44742-C4-1-R (CAVI-ROSE), and TIN2015-67522-C3-1-R (TRACES).

References

O. S. Bagge, K. T. Kalleberg, E. Visser, and M. Haveraaen. Design of the CodeBoost Transformation System for Domain-Specific Optimisation of C++ Programs. In *Third International Workshop on Source Code Analysis and Manipulation (SCAM 2003*, pages 65–75. IEEE, 2003.

I. D. Baxter, C. Pidgeon, and M. Mehlich. DMS®: Program transformations for practical scalable software evolution. In *Proceedings of the 26th International Conference on Software Engineering*, pages 625–634. IEEE Computer Society, 2004.

M. Boekhold, I. Karkowski, and H. Corporaal. Transforming and parallelizing ANSI C programs using pattern recognition. In *High-Performance Computing and Networking*, pages 673–682. Springer, 1999.

U. Bondhugula, A. Hartono, J. Ramanujam, and P. Sadayappan. A practical automatic polyhedral parallelizer and locality optimizer. *SIGPLAN Not.*, 43(6):101–113, June 2008. ISSN 0362-1340. doi: 10.1145/1379022.1375595. URL http://doi.acm.org/10.1145/1379022.1375595.

A. Brown, W. Luk, and P. Kelly. Optimising Transformations for Hardware Compilation. Technical report, Imperial College London, 2005.

J. R. Cordy. The TXL source transformation language. *Science of Computer Programming*, 61(3):190–210, 2006.

A. Danalis, G. Marin, C. McCurdy, J. S. Meredith, P. C. Roth, K. Spafford, V. Tipparaju, and J. S. Vetter. The Scalable Heterogeneous Computing (SHOC) Benchmark Suite. In *Proceedings of the 3rd Workshop on General-Purpose Computation on Graphics Processing Units*, pages 63–74. ACM, 2010.

C. Dave, H. Bae, S. Min, S. Lee, R. Eigenmann, and S. P. Midkiff. Cetus: A source-to-source compiler infrastructure for multicores. *IEEE Computer*, 42(11):36–42, 2009. doi: 10.1109/MC.2009.385. URL http://dx.doi.org/10.1109/MC.2009.385.

E. Di Napoli, D. Fabregat-Traver, G. Quintana-Orti, and P. Bientinesi. Towards an efficient use of the BLAS library for multilinear tensor contractions. *Applied Mathematics and Computation*, 235:454–468, May 2014.

D. Fabregat-Traver and P. Bientinesi. Application-tailored linear algebra algorithms: A search-based approach. *International Journal of High Performance Computing Applications (IJHPCA)*, 27(4):425 – 438, Nov. 2013.

F. Franchetti, Y. Voronenko, and M. Püschel. FFT Program Generation for Shared Memory: SMP and Multicore. In *Supercomputing (SC)*, 2006.

B. Huber. The Language.C Package. https://hackage.haskell.org/package/language-c, 2014.

B. Jacobs and F. Piessens. The VeriFast program verifier. Technical report, Department of Computer Science, KU Leuven, 2008.

B. Jacobs, J. Smans, P. Philippaerts, F. Vogels, W. Penninckx, and F. Piessens. Verifast: A powerful, sound, predictable, fast verifier for C and Java. In *Proceedings of the Third International Symposium on NASA Formal Methods, NFM 2011, Pasadena, CA, USA, April 18-20, 2011.*, pages 41–55, 2011. doi: 10.1007/978-3-642-20398-5_4. URL http://dx.doi.org/10.1007/978-3-642-20398-5_4.

ROCCC 2.0 User's Manual. Jacquard Computing Inc., revision 0.74 edition, June 2012. http://roccc.cs.ucr.edu/UserManual.pdf.

P. Klint, T. Van Der Storm, and J. Vinju. Rascal: A domain specific language for source code analysis and manipulation. In *Source Code Analysis and Manipulation, 2009. SCAM'09. Ninth IEEE International Working Conference on*, pages 168–177. IEEE, 2009.

R. Lammel, S. P. Jones, and J. P. Magalhaes. The SYB Package. https://hackage.haskell.org/package/syb, 2009.

O. Lindtjorn, R. G. Clapp, O. Pell, H. Fu, M. J. Flynn, and O. Mencer. Beyond Traditional Microprocessors for Geoscience High-Performance Computing Applications. *IEEE Micro*, 31(2):41–49, 2011.

Maxeler Technologies. Max Compiler MPT. https://www.maxeler.com/solutions/low-latency/maxcompilermpt/, March 2016.

P. W. O'Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In L. Fribourg, editor, *CSL*, volume 2142 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2001. ISBN 3-540-42554-3.

J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, pages 55–74. IEEE Computer Society, 2002. ISBN 0-7695-1483-9.

D. Rubio, C. W. Glass, J. Kuper, and R. de Groote. Introducing and exploiting hierarchical structural information. In *IEEE International Conference on Cluster Computing (CLUSTER), 2015*, pages 777–784. IEEE, 2015.

S. Schupp, D. Gregor, D. Musser, and S.-M. Liu. Semantic and behavioral library transformations. *Information and Software Technology*, 44(13): 797–810, 2002.

J. van Wijngaarden and E. Visser. Program Transformation Mechanics. Technical report, Technical Report UU-CS-2003-048, Universiteit Utrecht, 2003.

S. Verdoolaege and T. Grosser. Polyhedral extraction tool. In *Second International Workshop on Polyhedral Compilation Techniques (IMPACT'12), Paris, France*, 2012.

G. Vigueras, M. Carro, S. Tamarit, and J. Mariño. Towards Automatic Learning of Heuristics for Mechanical Transformations of Procedural Code. In *First Workshop on Program Transformation for Programmability in Heterogeneous Architectures (PROHA'16)*, March 2016.

E. Visser. Program Transformation with Stratego/XT: Rules, Strategies, Tools, and Systems in StrategoXT-0.9. In C. Lengauer, D. Batory, C. Consel, and M. Odersk, editors, *Domain-Specific Program Generation*, volume 3016 of *Lecture Notes in Computer Science*, pages 216–238. Springer-Verlag, June 2004.

E. Visser. A Survey of Strategies in Rule-Based Program Transformation Systems. *Journal of Symbolic Computation*, 40(1):831–873, 2005. Special issue on Reduction Strategies in Rewriting and Programming.

K. Willadsen. Meld. http://meldmerge.org/, February 2016. Retrieved on February 2016.